ERDC MSRC

Major Shared Resource Center

ERDC MSRC/PET TR/00-36

Analysis of HPC Usage: ERDC MSRC

by

Daniel Duffy

8 February 2001

02h00401

Work funded by the DoD High Performance Computing
Modernization Program ERDC
Major Shared Resource Center through

Programming Environment and Training (PET)

Supported by Contract Number: DAHC94-96-C0002
CSC Nichols

Analysis of HPC Usage: ERDC MSRC

Daniel Duffy*, William Ward*, Mark Fahey*, Rebecca Fahey*, and Tom Oppe†
Computational Migration Group
U.S. Army Engineer Research and Development Center
Major Shared Resource Center (ERDC MSRC)
3909 Halls Ferry Road
Vicksburg, MS 39180-6199

February 8, 2001

Abstract

As the Department of Defense prepares its newest version of a benchmark for vendors, an updated usage of computer platforms at all the Major Shared Resource Centers (MSRC) must be analyzed. This usage data provides the developers of the benchmark with information about the heaviest users of each machine and, hence, identifies which applications should be included. This report presents and analyzes the usage data for all four High Performance Computers (HPC) at the U.S. Army Engineer Research and Development Center (ERDC). Results show that the top five users of each machine accurately describe the total machine use. Furthermore, a summary of all four machines at ERDC will be given and a typical profile of their use presented. It is argued that any throughput benchmark test which uses a mix of codes should conform to this profile in order to accurately describe how the hardware is being used at ERDC.

1 Introduction

Originally established in 1989 as an Army Supercomputer Center, the U.S. Army Engineer Research and Development Center (ERDC) in Vicksburg, Mississippi, has developed into one of the top computing facilities in the world. ERDC performs necessary and critical research in the areas of Civil Engineering, Environmental Quality, and Environmental Sciences. In 1993 as part of the Department of Defense (DoD) High Performance Computing (HPC) Modernization Program, ERDC established the first HPC Major Shared Resource Center (MSRC).[1]

The ERDC MSRC operates multi-vendor HPC systems currently emphasizing large, scalable systems. Specifically, ERDC MSRC has four HPCs being used for DoD research. These include a 544 processor Cray T3E, a 128 processor SGI Origin 2000,[2] a 382 processor IBM SP complex (split into two images: 127 processors and 255 processors), and the IBM power3 symmetric multiprocessor with 64 distributed nodes, each with 8 processors. The users of these machines are categorized by computational technology areas (CTAs). The major CTAs using the hardware at ERDC MSRC include Computational Chemistry and Materials Science (CCM), Computational Electromagnetics and Acoustics (CEA), Computational Fluid Dynamics (CFD), Computational Electronics and Nanoelectronics (CEN), Computational Structural Mechanics (CSM), Environmental Quality Modeling and Simulation (EQM), Climate/Weather/Ocean Modeling and Simulation (CWO) and Signal/Image Processing (SIP).[3]

*Computer Sciences Corporation
†Data Management Consultants

[1] For more information about ERDC, visit the web site at http://www.wes.hpc.mil.

[2] As of December 1, 2000, the Origin 2000 was replaced by an Origin 3800 which will eventually have 512 processors. The data in this report only reflect the Origin 2000.

[3] For more information about the CTAs or the other three MSRCs, visit the High Performance Computing Modernization Program web site at http://www.hpcmo.hpc.mil.

Currently, the DoD is in the midst of developing a MSRC wide benchmark for future procurements. This benchmark will initially consist of applications currently being used at all four MSRCs. The choice of which applications to use and how best to run them will be made from an analysis of the usage data of each machine at all four MSRCs. The purpose of this report is to summarize and analyze the current usage of the machines at ERDC for the purpose of the DoD benchmark.

2 SGI Origin 2000

As stated above, the Origin 2000 has 128 processors available for computing. This machine is very user friendly and has the largest user base of all the machines at ERDC (approximately 145 different user names). It should be noted that the Origin is not currently being used to support the computationally intensive DoD challenge projects.[4]

Table 1(a) contains a summary of the usage data for the Origin 2000 based upon the log files created by the Portable Batch System (PBS) at ERDC MSRC. All user jobs from Jan. 1, 1999 to Sept. 8, 2000 have been counted in the tables below. To simplify, the runs were grouped into large bins of 16 processors each with time blocks of 2 node-hours.[5] Any run which took more than 6 node-hours was counted together. The table shows a percentage of usage summing up to 100% defined by the total number of node-hours in a given bin divided by the total number of node-hours used on the machine.

It is clear from the table that many jobs on the Origin require a range of processors from 1 up to 16. Of all the machines at ERDC, the Origin has the largest usage in this range of processors. Often, the Origin is used as a test platform for many codes. While testing, a code is generally run on a small number of processors. For production runs, though, the code usually requires a larger number of processors or may even be migrated to a different machine.

The largest usage of the Origin comes in the 6 hour and greater time range with the number of processors occurring between 17 and 32. Users tend to view the queue before submitting their jobs to decide on the number of processors to be used for that run. A run with up to 25% of the machine, i.e., up to 32 processors, will get through the queue faster than a run with a greater number of processors. Hence, the lower triangle in Table 1 is basically empty. Few jobs are submitted with a large number of processors which requires a relatively small amount of time. Users will sacrifice wall clock time to get their job in the queue and running rather than wait for a long time in the queue.

Table 1(b) contains the same data however only the top five users by total number of node-hours of the Origin are shown. Again, the percentages will sum up to be 100%; hence, the total number of node-hours in this table is defined to be the sum of the total number of node-hours used by only those top five users. On the Origin, these heavy users account for only 35% of the total usage of the machine. Even though the top five only account for about one-third of the total usage, their profile closely matches the profile for all users.

3 IBM Power3

The IBM Power3 symmetric multiprocessor (SMP) has 64 nodes with 8 processors each for a combined total of 512 processors. Due to hardware constraints, only 4 of the 8 processors on each node are available to be used for processes.[6] However, all 8 processors can be used if a code takes advantage of threads. The usage data, though, only accounts for processes and not threads;

[4]For more information about DoD Challenge Projects, visit the High Performance Computing Modernization Program web site at http://www.hpcmo.hpc.mil/.

[5]A node-hour is one hour per CPU. Hence, a job which requires 4 CPUs and runs for 2 wall clock hours gets charged a total of 8 node-hours.

[6]As of January 2001, the SwitchII had been installed, and all 8 processors on each node were available for both MPI and threaded jobs. The usage data used in this study only contains information prior to the installation of the SwitchII.

hence, the data shown below can only account for up to 256 processors. Furthermore, since the machine is new to the site, the log files on the IBM Power3 only go back to Feb. 22, 2000.

Table 2(a) shows the combined usage of 61 different user names. A significant percentage of use can be seen in the 1-16 processor range. However, the largest use of IBM Power3 occurs in the processor range of 113-128 processors which is around 25% of the total machine. Once the switch is installed to let users use the full 8 processors per node, the heaviest usage should occur around 50%, i.e., in the 256 processor range.

Table 2(b) shows the top five users of the Power3. For this machine, these five users account for 63.2% of the total machine's use. Therefore, it is fully expected that the profile of these users should closely match that of the total machine's use.

4 IBM SP

The IBM SP complex is split into two images: the first image is configured with 255 processors (Image A) and while the second has 127 processors (Image B). Table 3(a) and 4(a) show the total usage accumulated from 101 users on Image A and 70 users on Image B. A clear peak in the usage of Image A occurs at the 50% level or 128 processor range. Image B also has a peak at the same percentage of processors and exhibits large usage around the 75% node level.

Tables 3(b) and 4(b) show the usage for the top five users of Images A and B respectively, accounting for 65.5% of the total usage of Image A and 66.0% of the total usage of Image B. Again, the correlations are easily seen by comparing the tables.

5 T3E

The 544 processor T3E is used for some of the largest jobs, in terms of numbers of CPUs, currently being run at ERDC. Of the 544 processors, 520 are available for computations, and the maximum size for a single job is limited to 512 total CPUs. Table 5(a) shows the break down of usage for approximately 115 users. Peaks are seen at the 25%, 50%, and 75% levels of the machine, i.e., around 128, 256, and 384 processors.

The same behavior is seen in Table 5(b) which only accounts for the top five users of the system. These five users account for 53.1% of the total machine usage and have accumulated over 1.5 million node-hours, that is around 170 years of total computation time. Again, the correlations between the total usage and the top five users are obvious.

6 Cumulative Usage

The purpose of this exercise is to understand how the machines at ERDC are typically utilized. Therefore, any benchmark test used to describe their peformance should match the above profiles.[7] Hence, when running a throughput test, i.e., a mix of jobs run through a batch scheduler in order to describe a typical usage of an HPC, the numbers of jobs, the amount of time of each job, and the number of node-hours used should correlate to the above data. If the mix does not fit the profile of current usage, the results of the benchmark will not accurately represent the current usage of the HPCs at ERDC.

In Fig. 1, we show a complete histogram of the total amount of node-hours versus the number of CPUs used per job for all four machines at ERDC. To obtain the histogram, the number of node-hours for each job was placed into bins where the width of the bin for Fig. 1(a) is 1% of the total number of CPUs on any given machine. Note that the number of CPUs are different for each machine. Hence, the 1% bin on the Origin only represents those jobs which used 1 CPU. However, on the T3E, all jobs up to 7 CPUs are put into the 1% bin. After the usage across all machines was placed into these bins, the bin total was then divided by the total number of

[7] This idea is not unique and has been used by the developers of the *Effective System Performance* (ESP) benchmark at the University of California, Berkeley. More information can be found at http://www.nersc.gov/aboutnersc/esp.html.

node-hours used across all platforms in order to obtain a percentage. Thus, the area under the histogram integrates to 100%.

As can be seen in the figure, several large features occur. Namely, peaks occur at 12.5%, 25%, 50%, 75% and 100% CPUs used. Furthermore, there are a large number of jobs which request 0% up to 25% of the total machine. This makes sense simply because users are required to submit their jobs through a batch scheduler. Therefore, jobs which require a small number of CPUs can run concurrently with a job running on a large number of CPUs. Many users will use fewer numbers of CPUs to get their job running faster, even if that means running for an overall longer time.

In Fig. 1(b), the histogram is grouped into bin sizes of 5%. This tends to average out the smaller structures and give a more general usage of the machines at ERDC in order to see general trends. The peaks seen in Fig. 1(a) remain and have now been smoothed over. In general, the usage up to 25% remains high with a sharp drop off slightly above one-quarter of the available CPUs. The next sharp peak occurs at the 50% level followed by another quick drop. Slowly, the number of node-hours used begins to rise until another peak is seen around 75% of the available CPU's, and a final peak occurs at the machine limit. Any benchmark mix that tests throughput of the HPCs at ERDC should closely match this distribution.

As a comparison, Fig. 2 shows the same histogram computed for the top five users of each machine and summed across all platforms. The similarities between all users and only focusing on the top five heaviest users are again apparent.

7 Defining a Mix

As stated above, the usage data in Fig. 1 can be used to design a throughput test, or mix of programs, to benchmark the typical usage of the machines at ERDC. Assume that the machine to be tested has 512 processors and the total length of the test should not last more than 6 wall clock hours. Hence, the total amount of node-hours to use will be 3072 (6×512). Based on the percentage of node-hours used for each range of CPUs as shown in Fig. 1, Table 6 shows a set of runs which can be defined based on this distribution.

To further explain, jobs which require around 25% of the total number of CPUs on a machine account for around 14.7% of the total number of node-hours used at ERDC. Therefore, in a mix used to describe ERDC HPC usage, these types of jobs would account for 14.7% of the 3072 node-hours, or 450.8 node-hours. One way to break up these 450.8 node-hours is to run 3 runs of 1 hour each and 1 run of 31.3 minutes, each run using 32 CPUs.

In comparison, the runs which test the machine limits, i.e., those running at the maximum number of CPUs available, only account for 3.6% of ERDC's runs. Hence, 3.6% of the total 3072 node-hours contained in the mix would only be 110.3 node-hours. Thus, a single short run of 12.9 minutes using the maximum number of CPUs available would be sufficient to satisfy the profile at the machine limit.

Overall, as shown in Table 6, 33 separate jobs would be needed to test out the hypothetical 512 processor machine. Many of the jobs would be designed to only last 1 hour while others would last for shorter times. This is just one way in which the job mix may be defined. As long as the profile seen in Fig. 1 was conserved, many other job mixes may be defined depending on the requirements and time limits of the benchmark test; each job mix would be equally valid for ERDC.

Unfortunately, this data cannot be used to determine the order of the submission for all the various jobs defined for a mix. To accurately describe how jobs are submitted to the queue, certain assumptions will have to be made. A completely random submission of the jobs to the queue is not representative of how users submit jobs. Often, users will change the number of CPUs used for a job based on the number of CPUs free or based on the number of jobs in the queue. Hence, any order of submission must reflect some of these important issues.

8 Conclusions

In conclusion, the overall usage data for the four HPCs at ERDC was presented and analyzed. It was shown that the top five users of each machine may account for as much as 60% or more of a given machine. Hence, these five users dictate how each machine is used.

Combining the usage data into a single histogram reveals a profile of how the HPCs are utilized at ERDC. Some general trends can be extracted from the overall profile. That is, jobs tend to request blocks of CPUs around 25%, 50%, 75% and 100% of the total number of CPUs available. These correspond to jobs of size 32, 64, 96 and 128 CPUs on the SGI Origin 2000. Furthermore, many jobs fill in the gaps of each machine and run on less than 25% of the total number of CPUs. In many ways, this is a very efficient way to keep the usage of a machine close to 100%. In other words, if a user logs on to a machine and observes that only a few CPUs are idle, often that user will only request those CPUs which are not being used.

One possible use of this profile is to define the number, size, and time limit of jobs to be used as a benchmark mix. For a throughput measure of a platform that represents ERDC, jobs need to be run such that the profile fits that shown in Fig. 1. A profile which does not reflect the usage at ERDC will not represent how the hardware is currently being employed by the users.

In order to define the most representative benchmark mix for each MSRC, analysis of the usage data specific to each MSRC is necessary. Even so, since ERDC specializes in massively parallel machines, it is expected that other large parallel architectures would show similar profiles.

9 Acknowledgments

We acknowledge the help of Jay Cliburn, Mike Gough, and Owen LaGarde for helping to compile the usage data in this report.

10 Appendix: Tables and Figures

Table 1: Origin 2000: ERDC MSRC. NP is the range of processors used for each job and the range of times are broken into blocks of 2 hours each. All jobs over 6 hours are cumulated into the same bin.

(a) Cumulative usage of all users.

	Range of Times (hours)			
NP	0..2	3..4	4..6	6..
1-16	3.5	12.8	6.5	23.5
17-32	1.9	5.6	2.7	31.3
33-48	0.4	0.5	0.2	4.0
49-64	0.3	0.4	0.4	5.4
65-80	0.2	0.0	0.1	0.0
81-96	0.1	0.1	0.0	0.6
97-112	0.0	0.0	0.0	0.0
113-128	0.0	0.0	0.0	0.2

(b) Cumulative usage of the top 5 users only.

	Range of Times (hours)			
NP	0..2	3..4	4..6	6..
1-16	3.5	12.7	6.5	23.5
17-32	1.8	5.5	2.7	31.2
33-48	0.4	0.5	0.2	4.0
49-64	0.2	0.3	0.4	5.4
65-80	0.1	0.0	0.1	0.0
81-96	0.1	0.0	0.0	0.6
97-112	0.0	0.0	0.0	0.0
113-128	0.0	0.0	0.0	0.1

Table 2: IBM Power3: ERDC MSRC. Note that only 256 of the 512 can be currently accounted for under PBS.

(a) Cumulative usage of all users.

	Range of Times (hours)			
NP	0..2	3..4	4..6	6..
1-16	0.8	0.6	0.5	13.0
17-32	1.1	1.7	1.3	13.5
33-48	0.3	0.3	0.4	7.1
49-64	0.5	0.8	1.5	15.5
65-80	0.1	0.0	0.1	3.9
81-96	0.1	0.2	0.3	1.0
97-112	0.1	0.1	0.2	0.4
113-128	0.4	1.0	1.9	19.0
129-144	0.2	0.1	0.1	7.5
145-160	0.0	0.0	0.0	0.0
161-176	0.0	0.0	0.0	0.0
177-192	0.0	0.0	0.0	0.0
193-208	0.0	0.0	0.0	0.0
209-224	0.1	0.0	0.0	4.4
225-240	0.0	0.0	0.0	0.0
241-256	0.0	0.0	0.0	0.9

(b) Cumulative usage of the top 5 users only.

	Range of Times (hours)			
NP	0..2	3..4	4..6	6..
1-16	0.8	0.5	0.5	12.9
17-32	1.1	1.7	1.3	13.5
33-48	0.3	0.3	0.3	7.1
49-64	0.4	0.7	1.5	15.5
65-80	0.1	0.0	0.1	3.9
81-96	0.0	0.1	0.3	1.0
97-112	0.0	0.0	0.2	0.4
113-128	0.4	0.9	1.9	19.0
129-144	0.2	0.1	0.1	7.5
145-160	0.0	0.0	0.0	0.0
161-176	0.0	0.0	0.0	0.0
177-192	0.0	0.0	0.0	0.0
193-208	0.0	0.0	0.0	0.0
209-224	0.1	0.0	0.0	4.4
225-240	0.0	0.0	0.0	0.0
241-256	0.0	0.0	0.0	0.9

Table 3: IBM SP (Image A: 255 Processors): ERDC MSRC.

(a) Cumulative usage of all users.

	Range of Times (hours)			
NP	0..2	3..4	4..6	6..
1-16	1.2	1.5	2.2	3.0
17-32	2.7	2.8	5.6	7.8
33-48	0.3	1.5	2.9	3.1
49-64	0.2	0.2	0.3	5.4
65-80	0.1	0.0	0.1	0.0
81-96	0.1	0.2	0.1	1.2
97-112	0.1	0.0	0.0	3.6
113-128	0.5	0.5	0.4	20.6
129-144	0.0	0.0	0.0	1.3
145-160	0.1	0.2	0.0	9.1
161-176	0.1	0.0	0.0	5.1
177-192	0.0	0.0	0.2	2.2
193-208	0.2	0.2	0.6	4.5
209-224	0.3	0.0	0.0	7.0
225-240	0.1	0.1	0.0	2.0
241-255	0.0	0.0	0.0	0.0

(b) Cumulative usage of the top 5 users only.

	Range of Times (hours)			
NP	0..2	3..4	4..6	6..
1-16	1.2	1.5	2.1	3.0
17-32	2.7	2.8	5.5	7.7
33-48	0.3	1.5	2.9	3.1
49-64	0.2	0.2	0.2	5.4
65-80	0.0	0.0	0.1	0.0
81-96	0.1	0.2	0.1	1.2
97-112	0.0	0.0	0.0	3.6
113-128	0.4	0.4	0.3	20.5
129-144	0.0	0.0	0.0	1.3
145-160	0.1	0.2	0.0	9.0
161-176	0.0	0.0	0.0	5.1
177-192	0.0	0.0	0.2	2.2
193-208	0.1	0.2	0.6	4.4
209-224	0.3	0.0	0.0	7.0
225-240	0.0	0.1	0.0	2.0
241-255	0.0	0.0	0.0	0.0

Table 4: IBM SP (Image B: 127 Processors): ERDC MSRC.

(a) Cumulative usage of all users.

	Range of Times (hours)			
NP	0..2	3..4	4..6	6..
1-16	1.2	4.6	2.3	5.0
17-32	1.1	6.1	2.9	7.9
33-48	1.2	2.3	4.2	14.3
49-64	0.8	0.6	0.3	15.1
65-80	0.1	0.1	0.2	0.2
81-96	0.2	0.1	0.0	18.2
97-112	0.2	0.4	0.4	3.1
113-127	0.2	0.1	0.1	7.5

(b) Cumulative usage of the top 5 users only.

	Range of Times (hours)			
NP	0..2	3..4	4..6	6..
1-16	1.2	4.5	2.2	5.0
17-32	1.1	6.1	2.8	7.9
33-48	1.2	2.2	4.2	14.2
49-64	0.8	0.6	0.2	15.1
65-80	0.0	0.1	0.1	0.1
81-96	0.1	0.1	0.0	18.2
97-112	0.2	0.4	0.4	3.1
113-127	0.1	0.1	0.1	7.4

Table 5: T3E: ERDC MSRC.

(a) Cumulative usage of all users.

	Range of Times (hours)			
NP	0..2	3..4	4..6	6..
1-16	0.3	0.6	0.3	0.5
17-32	0.6	1.6	1.0	2.9
33-48	0.5	0.5	0.2	1.7
49-64	2.0	11.2	3.4	2.6
65-80	0.1	0.1	0.3	2.7
81-96	0.1	0.3	0.4	0.6
97-112	0.3	0.6	1.5	8.0
113-128	0.7	0.4	0.5	6.1
129-144	0.0	0.0	0.0	0.1
145-160	0.2	0.1	0.0	0.3
161-176	0.0	0.0	0.0	0.0
177-192	0.0	0.0	0.0	0.2
193-208	0.1	0.1	0.1	0.8
209-224	0.1	0.4	0.4	2.7
225-240	0.1	0.1	0.2	3.2
241-256	0.5	0.3	0.6	17.3
257-272	0.0	0.0	0.0	0.0
273-288	0.0	0.0	0.0	0.0
289-304	0.1	0.1	0.3	0.8
305-320	0.0	0.0	0.0	0.0
321-336	0.0	0.0	0.0	0.0
337-352	0.0	0.0	0.1	0.2
353-368	0.1	0.1	0.0	0.2
369-384	0.1	0.1	0.4	11.7
385-400	0.1	0.2	0.0	0.0
401-416	0.0	0.0	0.0	0.0
417-432	0.0	0.0	0.0	0.0
433-448	0.0	0.0	0.0	0.0
449-464	0.0	0.1	0.0	0.0
465-480	0.0	0.0	0.0	0.0
481-496	0.0	0.0	0.0	0.4
497-512	0.1	0.1	0.3	6.2

(b) Cumulative usage of the top 5 users only.

	Range of Times (hours)			
NP	0..2	3..4	4..6	6..
1-16	0.3	0.6	0.2	0.4
17-32	0.6	1.6	0.9	2.8
33-48	0.4	0.5	0.2	1.6
49-64	2.0	11.2	3.4	2.5
65-80	0.1	0.1	0.3	2.6
81-96	0.1	0.2	0.4	0.6
97-112	0.3	0.6	1.5	8.0
113-128	0.7	0.4	0.5	6.1
129-144	0.0	0.0	0.0	0.0
145-160	0.1	0.1	0.0	0.3
161-176	0.0	0.0	0.0	0.0
177-192	0.0	0.0	0.0	0.1
193-208	0.1	0.1	0.0	0.8
209-224	0.1	0.4	0.4	2.7
225-240	0.1	0.0	0.2	3.2
241-256	0.4	0.3	0.6	17.3
257-272	0.0	0.0	0.0	0.0
273-288	0.0	0.0	0.0	0.0
289-304	0.0	0.1	0.2	0.8
305-320	0.0	0.0	0.0	0.0
321-336	0.0	0.0	0.0	0.0
337-352	0.0	0.0	0.1	0.1
353-368	0.0	0.0	0.0	0.2
369-384	0.0	0.0	0.4	11.7
385-400	0.1	0.2	0.0	0.0
401-416	0.0	0.0	0.0	0.0
417-432	0.0	0.0	0.0	0.0
433-448	0.0	0.0	0.0	0.0
449-464	0.0	0.0	0.0	0.0
465-480	0.0	0.0	0.0	0.0
481-496	0.0	0.0	0.0	0.4
497-512	0.1	0.0	0.2	6.2

Table 6: Assuming a machine size of 512 processors for a 6 hour throughput test, the following table shows a possible mix of job sizes and times as defined by the profile of the typical usage of the machines. The types of runs are based on the maximum number of CPUs within the specified range of processors with the total number of runs being 33.

Range of CPUs	Node-hours	Percentage	Types of Runs
0..26	223.59	7.28	8 run(s) for 1 hour 1 run for 44.0 minutes
27..51	383.50	12.48	7 run(s) for 1 hour 1 run for 29.4 minutes
52..77	400.62	13.04	5 run(s) for 1 hour 1 run for 13.0 minutes
78..102	245.60	7.99	2 run(s) for 1 hour 1 run for 23.9 minutes
103..128	450.77	14.67	3 run(s) for 1 hour 1 run for 31.3 minutes
129..154	49.98	1.63	1 run for 19.5 minutes
155..179	16.05	0.52	1 run for 5.4 minutes
180..205	93.16	3.03	1 run for 27.3 minutes
206..230	84.53	2.75	1 run for 22.0 minutes
231..256	522.89	17.02	2 run(s) for 1 hour 1 run for 2.6 minutes
257..282	7.66	0.25	1 run for 1.6 minutes
283..307	19.13	0.62	1 run for 3.7 minutes
308..333	52.51	1.71	1 run for 9.5 minutes
334..358	95.80	3.12	1 run for 16.0 minutes
359..384	194.27	6.32	1 run for 30.4 minutes
385..410	43.13	1.40	1 run for 6.3 minutes
411..435	39.70	1.29	1 run for 5.5 minutes
436..461	0.70	0.02	1 run for 0.1 minutes
462..486	32.95	1.07	1 run for 4.1 minutes
487..512	110.33	3.59	1 run for 12.9 minutes

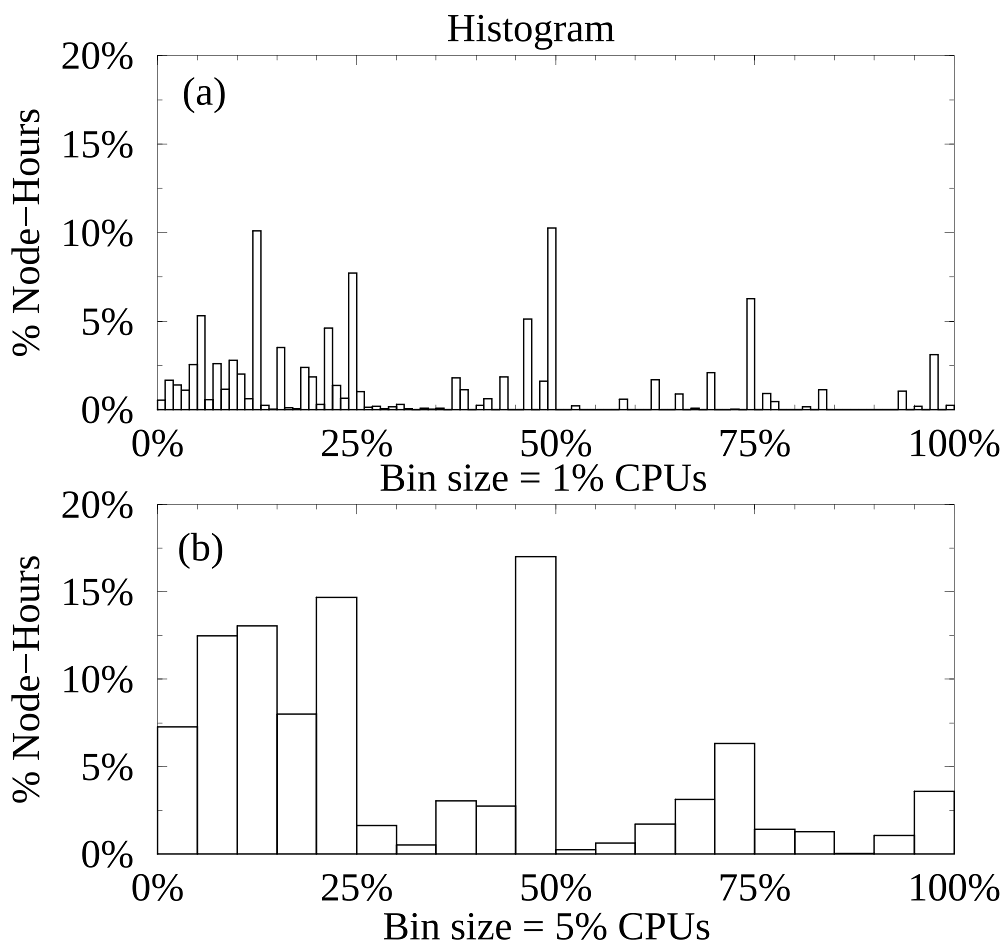

Figure 1: A histogram of the percentage of node-hours used per percentage of CPUs used per job on all four machines at ERDC. In (a), each bin represents 1% of CPUs used per machine. In (b), each bin represents 5% of the total number of CPUs used on each machine.

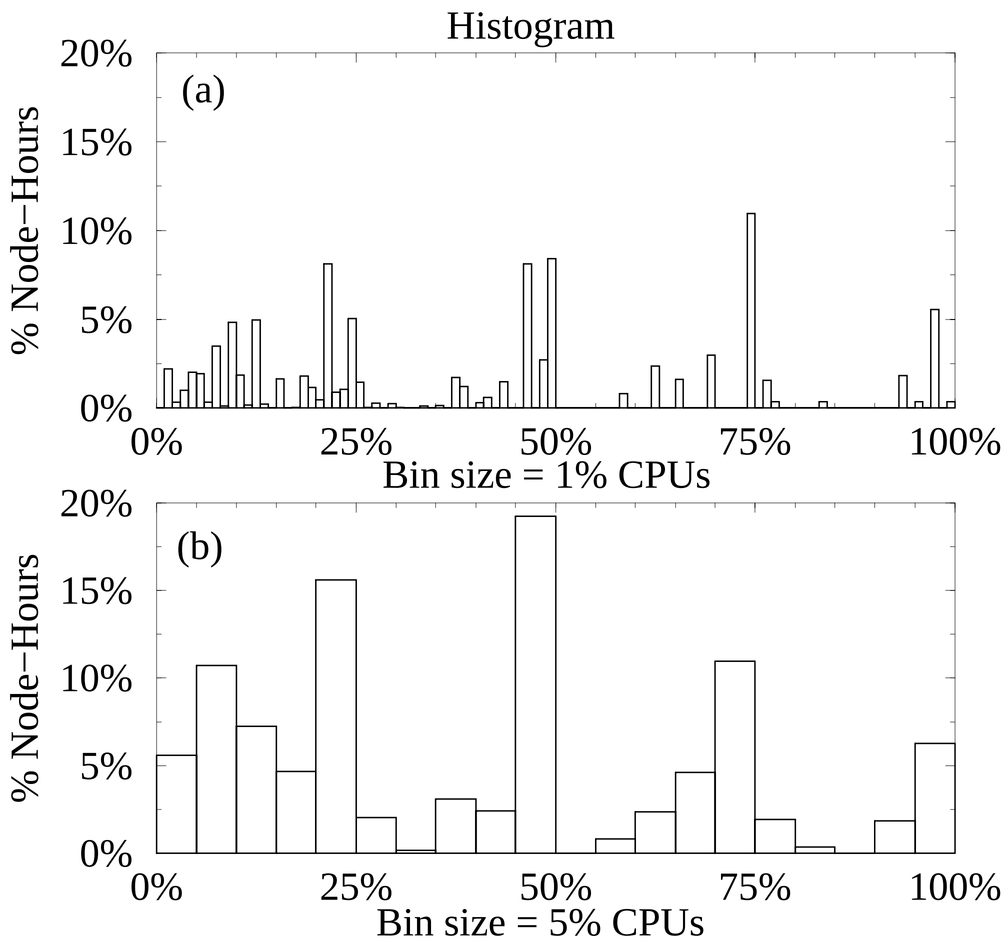

Figure 2: A histogram of the percentage of node-hours used per percentage of CPUs used per job on all four machines at ERDC for the top 5 users of each machine only. In (a), each bin represents 1% of CPUs used per machine. In (b), each bin represents 5% of the total number of CPUs used on each machine.